


Introduction

- Hawkes Processes are widely used to learn influence from event sequences



- Characterized by its intensity function

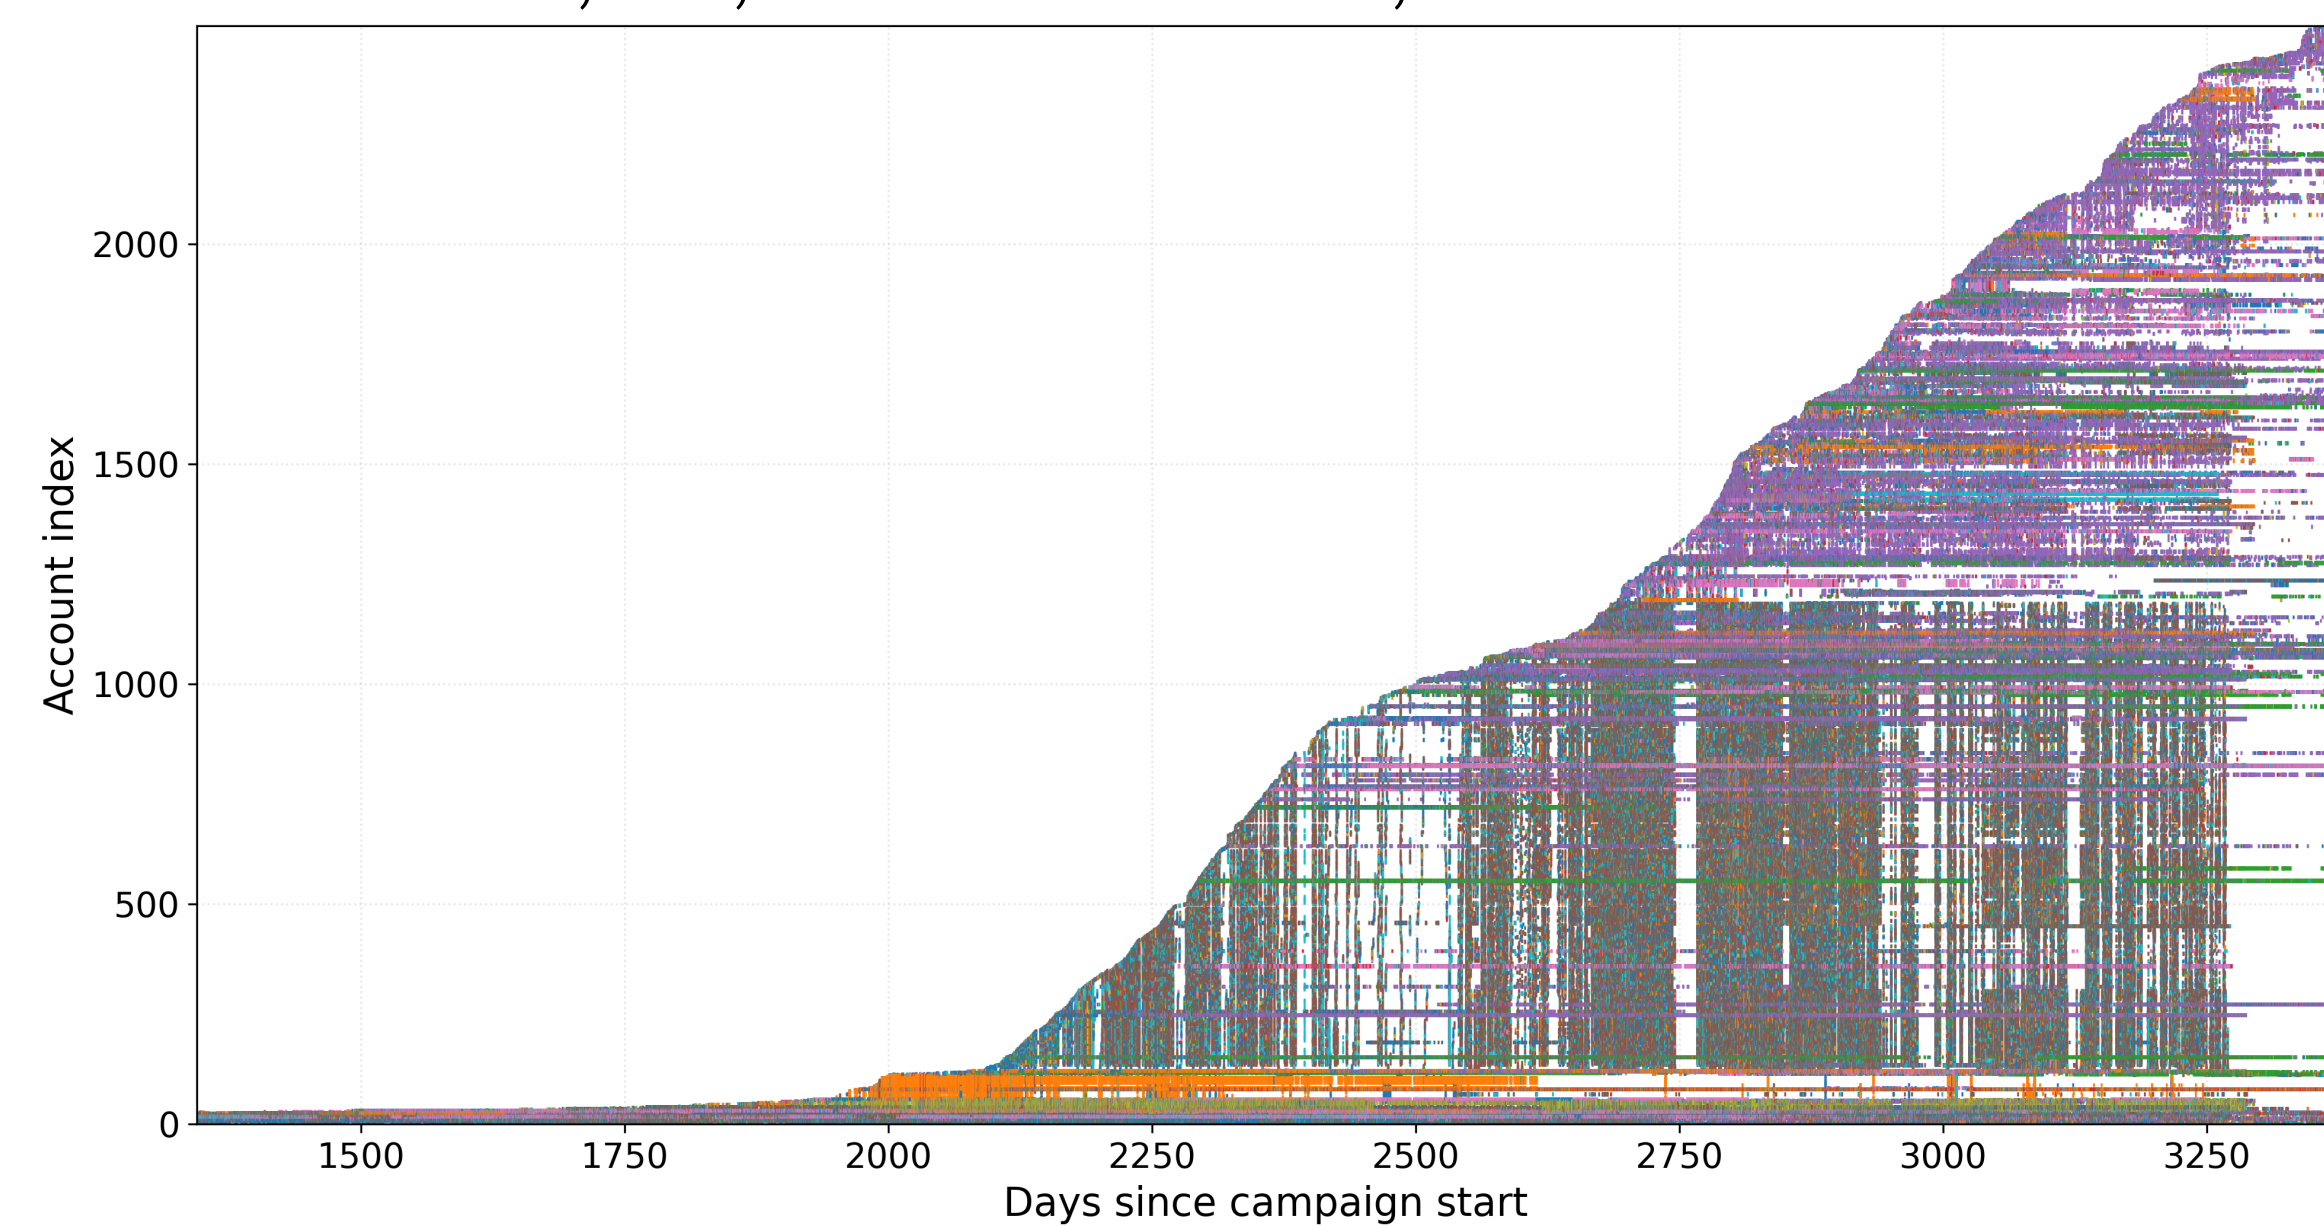
$$\lambda_p(t) = \mu_p + \sum_{j:t_j < t} \sum_{k=1}^K \alpha_{p,m_j,k} \gamma_k e^{-\gamma_k(t-t_j)}$$

- Issue: Fitting is $O(N^2)$ naive, $O(N)$ with recurrence, but unparallelizable

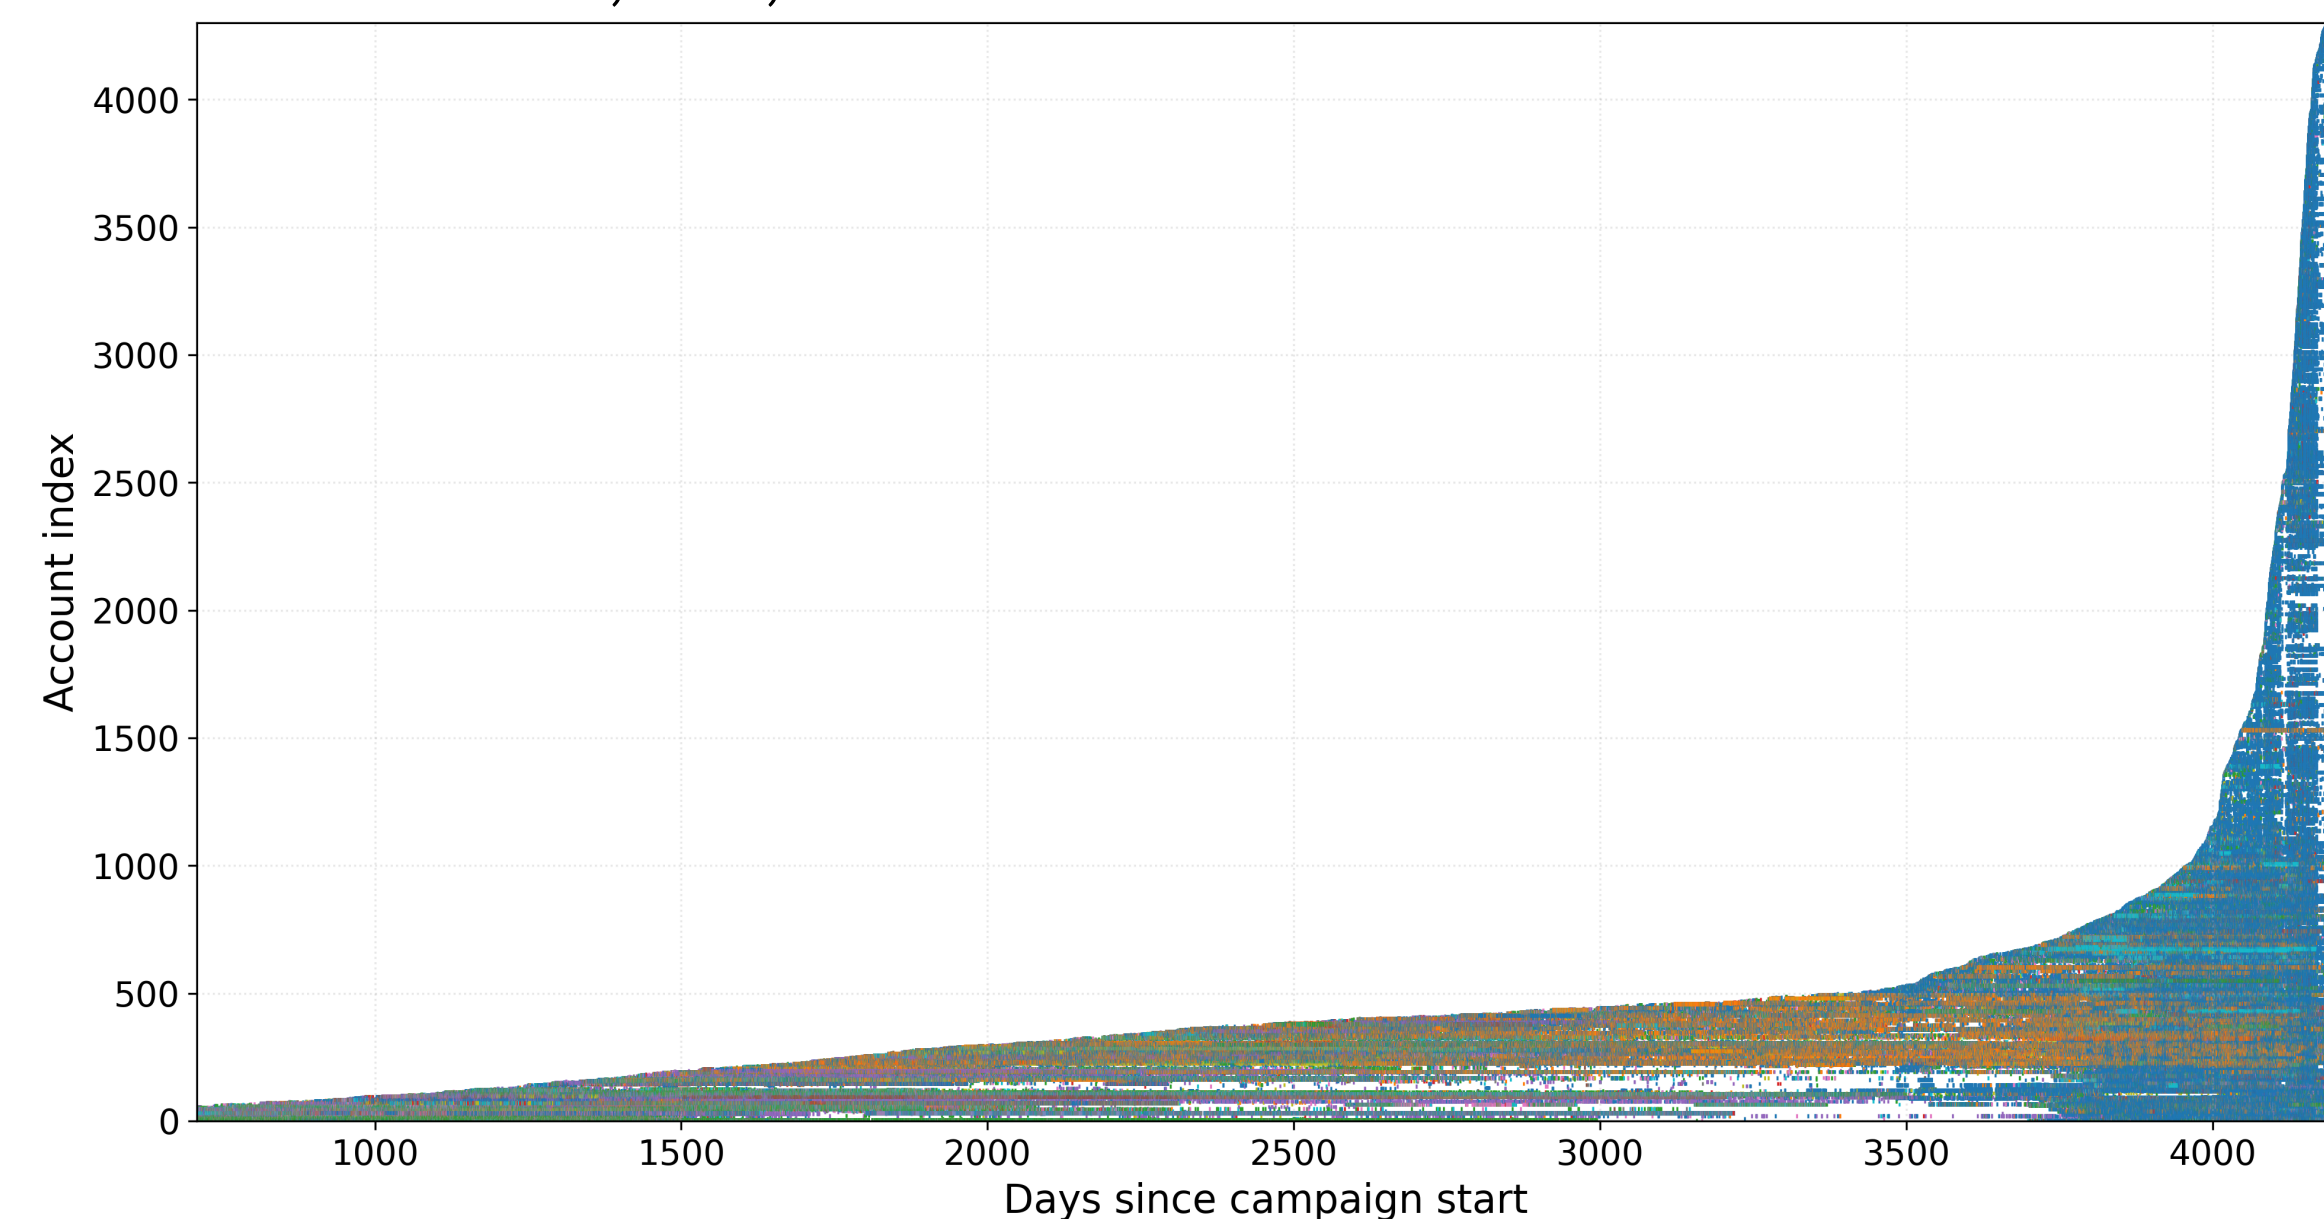
Motivating Application

- Twitter Information Operation Archive dataset
- 217M tweets from 87K accounts
- Linked to 47 state-backed operations
- How do these accounts influence each other?
- This scale of data is intractable for Hawkes processes

4,671,959 tweets from 2,496 accounts




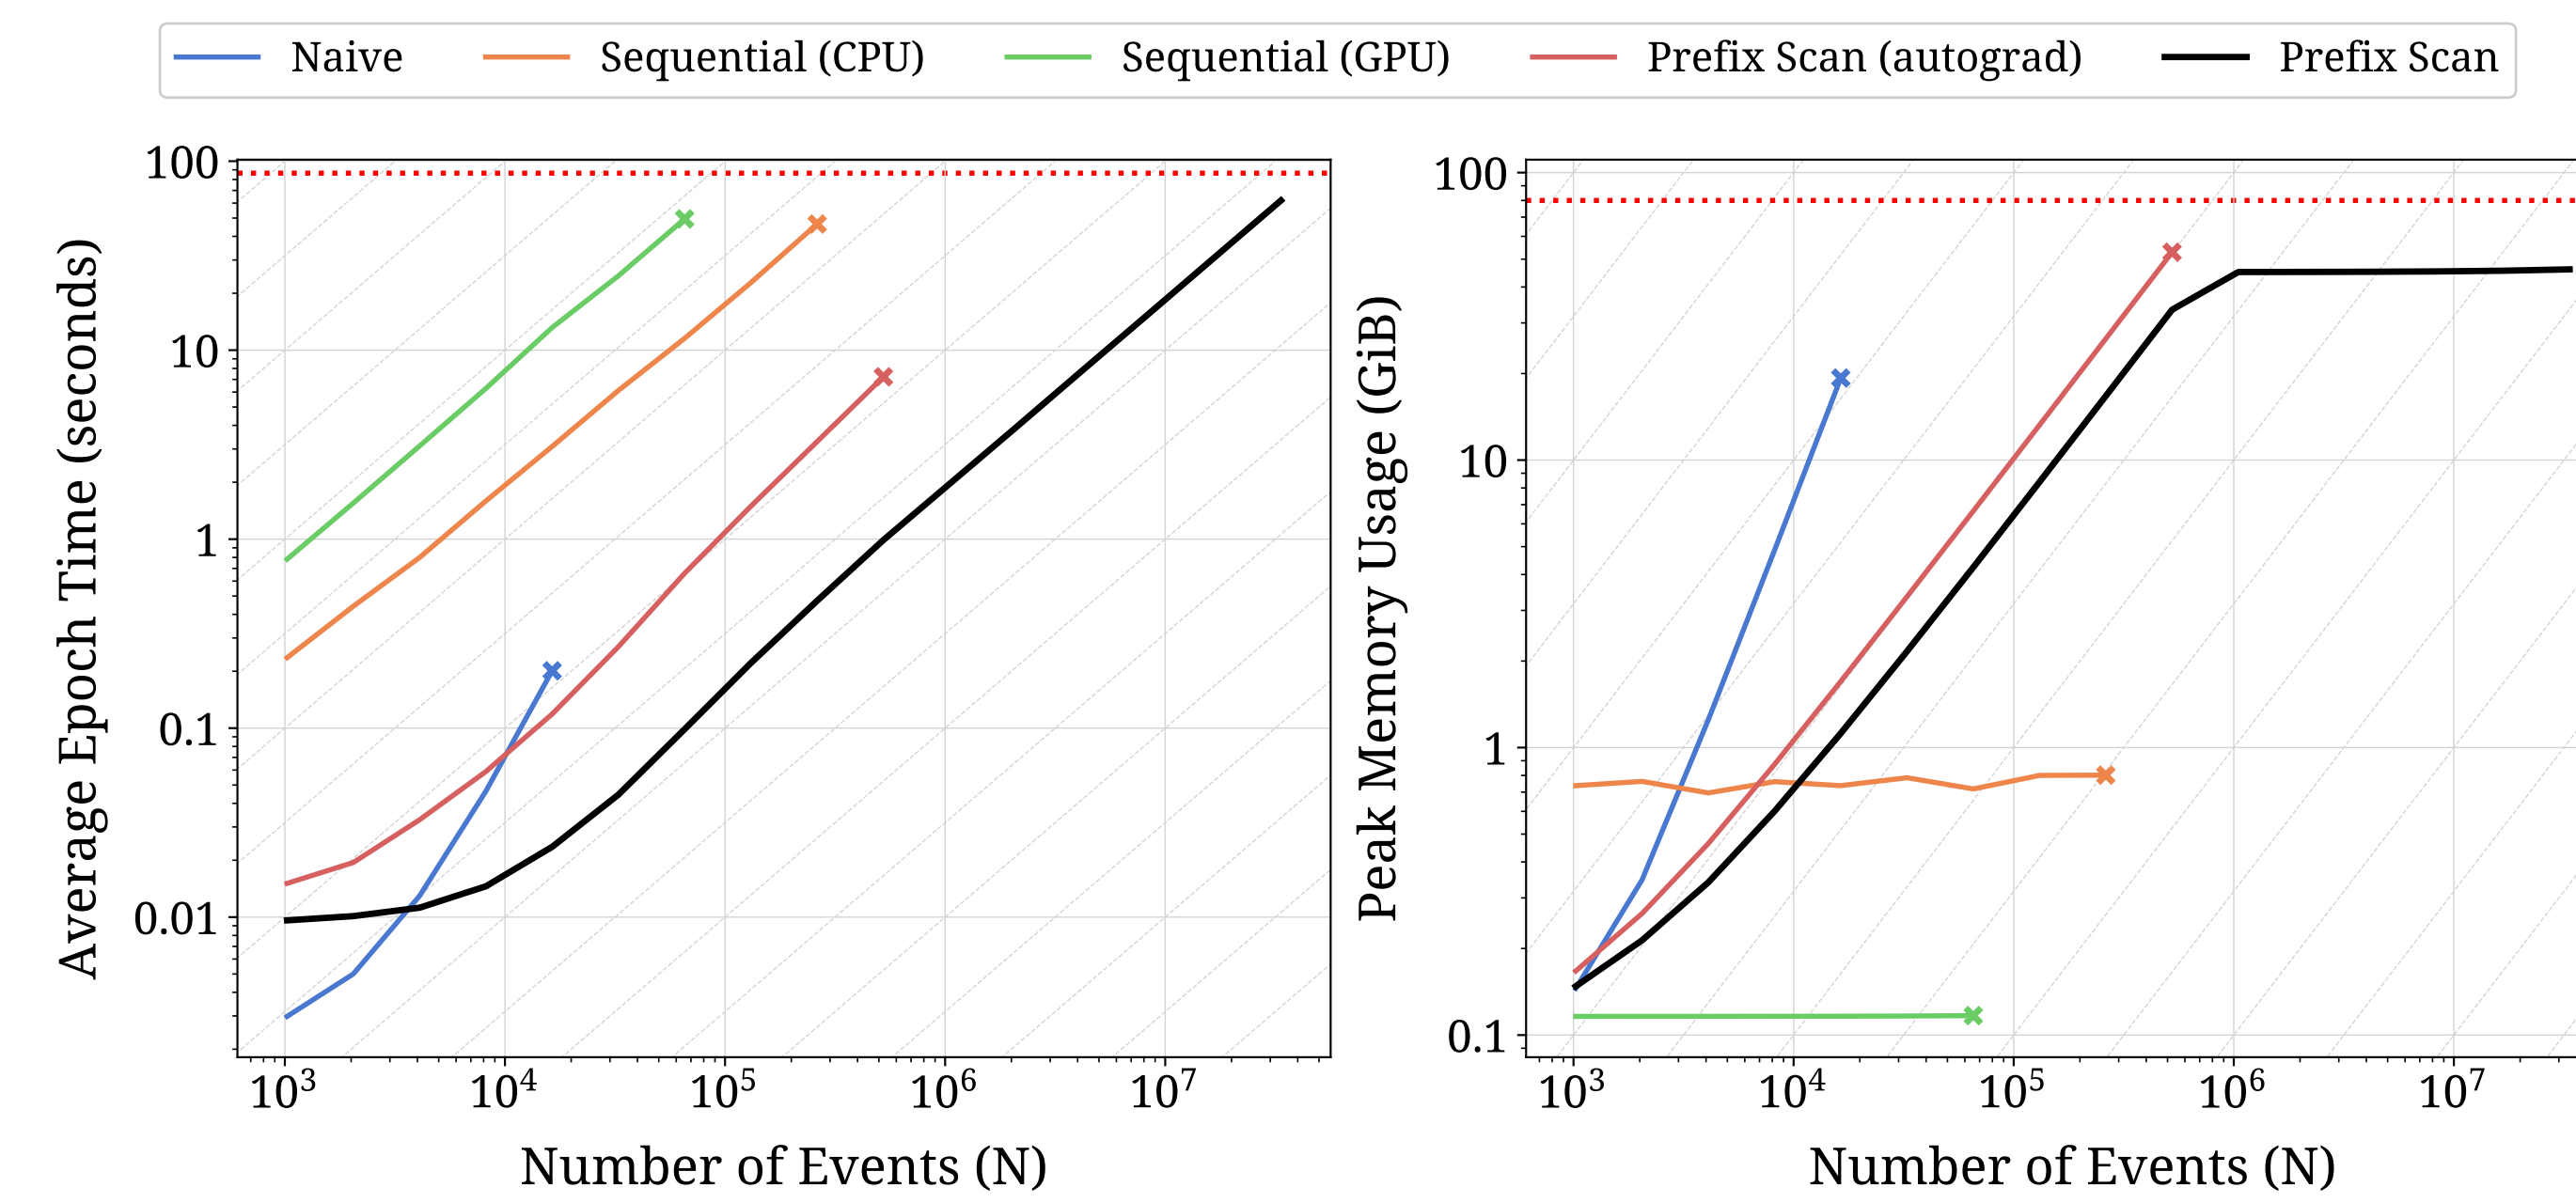
10,241,542 tweets from 4301 accounts



Our Contribution

A massively parallel algorithm for Hawkes modeling

- Speeds up estimation by orders of magnitude
- Constant memory usage for GPU computation
- No additional assumptions or approximations!



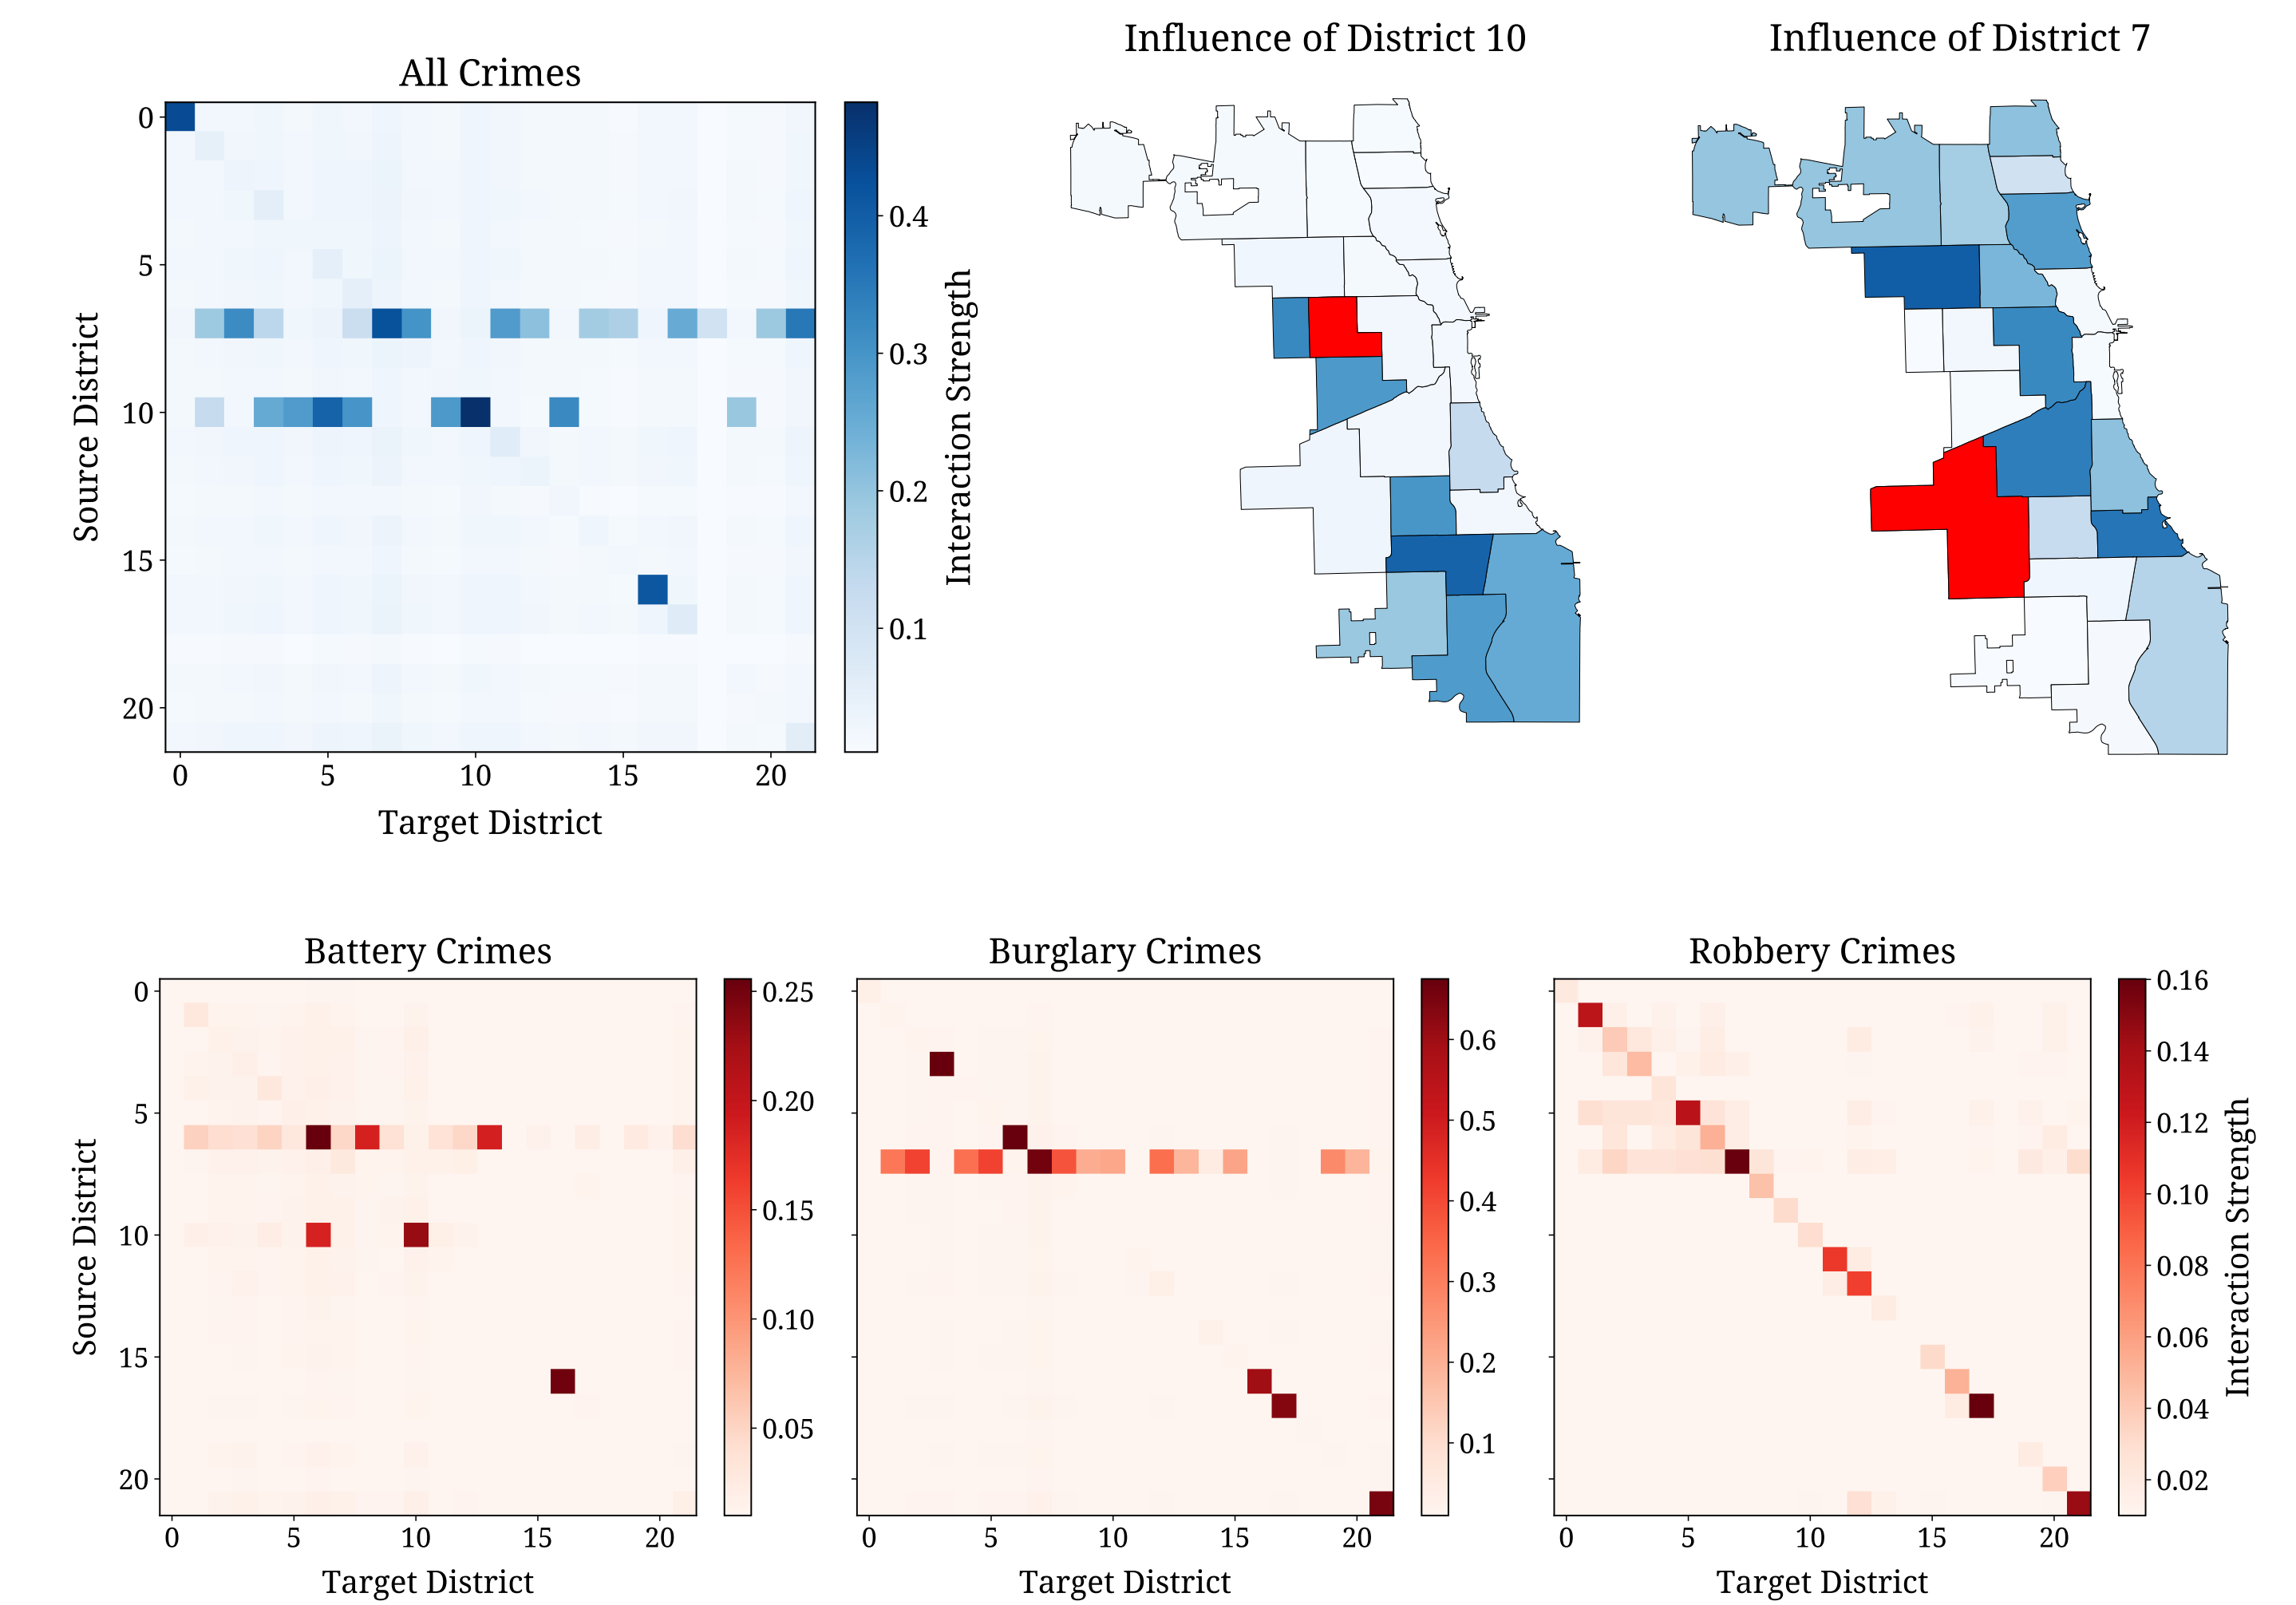
MemeTracker

- 34,547,727 unique memes posted on 5000 most active websites
- First time the Hawkes process has been applied to data of this scale

	Subset	Nodes	Events	Log Lik.	Avg. Epoch Time	
					Ours	Lazy MHP
	Bail Out	835	64,138	3.612	0.06 s	0.47 s
	Miami Heat	1,173	133,451	1.680	0.21 s	1.60 s
	Amy Winehouse	1,561	226,247	2.043	0.34 s	3.30 s
	Arab Spring	1,377	400,199	2.796	0.46 s	4.60 s
	Greece	1,925	1,077,125	2.280	1.24 s	—
	Libya	2,003	2,486,030	2.679	2.97 s	—
	Crisis	2,266	5,085,759	2.495	6.94 s	—
	Leader	2,344	10,604,129	2.674	14.91 s	—
	North Korea...	2,451	16,408,957	2.747	23.73 s	—
	Prince William...	2,494	32,687,780	2.827	48.33 s	—

Chicago Crime Data

- Chicago Police Department (CPD) provides all crimes committed since 2001
- 8,448,627 total crimes across 22 districts



COVID-19 Policy Diffusion

How did countries influence each other in implementing COVID-19 policies?

